REMARKS

Claims 1-34 are pending in the application. With entry of the amendment, claims 17, 18 and 30-34 are canceled and claims 1, 4, 11-13, 16 and 19 are amended.

Claim 11 was rejected under 35 USC 112, second paragraph. In particular, the phrase "said plurality" was objected to for lacking antecedent basis. Accordingly, claim 11 has been amended to address the rejection. Claim 16 has been amended in a similar manner. Reconsideration of the rejection under 35 USC 112 is respectfully requested.

Claims 1-4, 10, 12-15, 31 and 34 were rejected under 35 USC 102(e) as being anticipated by Uzun, et al. (US6961342)("Uzun"). The rejection is respectfully traversed.

Uzun is directed to uniform traffic, i.e., traffic that has no statistical variations, and it reconfigures its central switching fabric once per time slot.

The subject invention is founded upon the recognition that there is a degree of uniformity in traffic but there are also statistically significant variations between, for example, different states of uniform traffic.

Claim 1 has been amended to clarify the invention. In particular, claim 1 now recites in part "a packet switch for switching packets, carried in time slots, between traffic sources and traffic sinks, ... to allow transfer of packets according to a matrix with elements that specify the time-average required capacity between each traffic source – traffic sink pair, wherein the set of links is varied in accordance with variations in the time-average capacity and not at a time slot rate."

Conventional switch fabrics are reconfigured on a packet-by-packet basis. To do so requires the calculation every timeslot of a matching between input ports and output ports, and then switches between input port and output ports to be "made" to allow the correct connections. The calculation requires queue state information to be communicated from the input blocks (and possibly the output blocks) to a centralized scheduler that must complete its calculation within a timeslot measured in nanoseconds. In such a short timescale, it is not possible to calculate matchings that optimize throughput and delay and recourse is made to ad-hoc heuristics. Bottlenecks in state-of-the-art packet switches may be traced to the speed limitations of the centralized scheduler.

Consider an NxN route-and-select crossbar switch equipped with cross-point queues. Each cross-point queue is split into a queue tail and a queue head. The queue-tails are located with the route switch within an input block and the queue-heads are located with the select switch while preserving the one-to-one interconnection between the tails and heads of the corresponding cross-point queues. This leads to N input blocks and N output blocks. Of these, the N input blocks each have a single entrance port and N exit ports containing N cross-point queue tails; and the N output blocks each have a single exit port and N entrance ports containing N cross-point queue heads.

There is a transpose interconnect between the input and output blocks. The transpose interconnect provides full-mesh connectivity between the input and output blocks at the full-line rate, i.e., as many as N^2 packets can traverse the inter-stage interconnect in one timeslot. However, at most N packets can enter the switch in one timeslot. The interconnect therefore has a speed-up of N and as such is heavily over-engineered.

From a throughput perspective what is required is full-mesh connectivity between blocks, not necessarily at the full line rate, but instead at a rate that exceeds the bandwidth demand on each link (i.e., the bandwidth required by the flow between the input and output port pair associated to the specified cross-point).

This may be achieved by time division multiplexing (TDM). That is, the N exit ports of the input blocks are replaced by a single TDM exit port; the N entrance ports of the output ports are replaced by a single TDM entrance port; and the exit ports of the input blocks and the entrance ports of the output blocks are interconnected via a NxN space switch that cycles through a periodic schedule of permutations. The permutations required are calculated using a suitable algorithm to decompose a "service matrix" that majorizes a "traffic matrix" (or demand matrix). This ensures that the available bandwidth on any link exceeds the demand and hence that queue backlogs do not build up. The number of permutations required in each "frame" (period of the schedule) is finite.

This frame of linking operations will remain in operation without build up of queues or other problems while the traffic matrix/demand matrix pertains. When the traffic matrix varies on a macro level, i.e., the demand changes at the statistical level, the linking operations that

make up the frame must change to cope with the variation. In some embodiments variation of the traffic matrix is sensed by change in one or more queues

In an embodiment of the switch a single space switch is substituted by a collection of space switches placed between a set of spatial de-multiplexers and multiplexers: one associated with each input block and each output block respectively (see for example Fig. 29).

Then, the de-multiplexers and multiplexers simply select their output and input ports in a periodic cycle, thus selecting an individual space switch in each slot. The space switches need now only be reconfigured at most once every cycle. Essentially, the original frame of permutations has been laid out in two dimensions: one dimension in space the other in time. If the one space dimension exhausts the permutations within a frame, the space switches are no longer required to change configuration in time. Reconfiguration is only required to implement a new frame of permutations required to adapt to a change in traffic matrix.

Hence the switch architecture enables reconfiguration on a slower timescale than the "slot time" that characterizes the packets (e.g., the duration of a cell at the line rate).

Indeed, if statistically stationary traffic were experienced and an architecture with sufficient spatial speed-up, the central switch fabric need be configured once only.

In the case of non-stationary traffic, switch re-configuration is necessary to adapt to changing traffic patters only on a timescale that characterizes the non-stationary properties of the traffic.

Networks use traffic engineering processes to control the peak traffic on long time scales, whereas for short times scales, queuing processes iron out traffic fluctuations. Thus, in practice, switch fabric reconfiguration is therefore only required on time scales intermediate between these short term and long term smoothing mechanisms.

Again, it is possible to engineer the network so that reconfiguration is needed only rarely. The requirement for infrequent reconfiguration again facilitates the implementation of the central switch fabric in optical technology – current transparent switch technology scalable to large port counts tend to limit reconfiguration times to microsecond time scales. Sub-nanosecond reconfiguration is required for the packet-by-packet switching in opaque implementations. The method does not however restrict the implementation to optical technology.

The architecture dispenses with the need for a real-time centralized switch fabric scheduler. A centralized scheduler IS however required to calculate the switch configuration but the calculation may proceed essentially "off-line" as it only requires infrequent (in terms of timeslots) updates from the input blocks on traffic statistics. All high speed (packet by packet) scheduling occurs only locally within the blocks.

These features permit the implementation of switch fabrics using essentially passive and near zero energy consumption technology. This contrasts with the energy hungry opaque electronic switch fabrics with centralized schedulers of other switch architectures. Indeed "business as usual" in packet switch architectures is unsustainable. Scaled by current traffic predictions, or alternatively scaled according to the goal of a good proportion of the world's population enjoying at least 100 Mbits of broadband access, the network will consume more than the global electricity generating capacity and core internet routers will have operating temperatures comparable to the surface temperature of the sun. New solutions are therefore required. While there remains further work to be done on reducing energy consumption of the blocks, the invention contributes to these new solutions by eliminating the energy consumption of the switch fabric which can be many kilowatts.

Uzun describes a switch consisting of a central switch fabric sandwiched between a first stage of input blocks and a third stage of output blocks. Each input-block serves a group of input ports and an equal or larger group of M entrance ports to the switch fabric. Each output-block serves a group of output ports and an equal or large group of M exit ports from the switch fabric.

The grouping in Uzun permits up to M packets to depart from an input block each time slot even if there are less than M queues occupied within the input block. Similarly, it permits up to M packets to arrive at an output-block in each timeslot even if there are less than M output ports to the block. The queues within the blocks may be organized as Virtual Output Queues or Grouped Virtual Output Queues within the input blocks and as straight Output Queues within the output block.

The central switch fabric is reconfigured every slot and the switching device of the fabric switches at the bit rate. Although in principle it might be possible to modify Uzun to instead use a fabric that is transparent (i.e., no constraint is placed by the switches on the signals carried), the requirement for rapid reconfiguration would result in great expense.

The focus of Uzun's disclosure is on the scheduling algorithm. The switch structure permits the scheduling to be conceptually separated into finding matchings between groups and subsequently scheduling the links to maximize throughput (by minimizing the number of idle links in each slot).

To avoid the problems associated with real-time centralized scheduling described in the preceding, a simple cyclic permutation is chosen for the matchings between groups. This is restricted to the special case of the decomposition of a uniform service matrix. Uzun does not disclose nor envisage the possibility of adapting the schedule of matchings to better serve a non-uniform inter-group traffic matrix.

Rather, the scheduling of the grouped links is adjusted via ad hoc methods to improve performance in the presence of non-uniform traffic while attempting to keep the state-information required essentially localized to the input block.

Three variations are described in Uzun of increasing scheduler sophistication.

The simplest variation attempts to maximize the utilization of the edges (composed of a group of M links) of the matchings between groups. Grouping is fundamental to this process because it permits several packets destined for the same output port to depart from an input group when there are less than M packets pending with distinct output port destinations within the same group. Nevertheless, in each time-slot, it may not be possible to utilize all M links because there may be fewer than M packets destined for the output group pending at the input group. There will then be idle links between the switch fabric and the output group in that time-slot.

The second variation attempts to make use of the idle links by making them available for reservation by other input blocks. Link reservation first proceeds according to the simple scheme, idle links are then advertised for possible reservation to the other input blocks by a token ring. Reservations are released every time-slot.

The third variation removes the requirement for immediate release of link reservations by providing another token ring to handle the release mechanism. This provides a dynamic mechanism to adapt to non-uniform fluctuating traffic demands.

While it is understood that Uzun limits the formal computational complexity of the scheduling algorithm, the first algorithm performs poorly for non-uniform traffic and the second

and third algorithms are complex; the signaling (token rings) is not localized; and scheduling and signaling must all be performed within a timeslot i.e., in real time.

From the foregoing discussion, it is clear that Uzun does not teach "a packet switch <u>for switching packets, carried in time slots, between traffic sources and traffic sinks, ..., to allow transfer of packets according to a matrix with elements that specify the time-average required capacity between each traffic source – traffic sink pair, wherein the set of links is varied in accordance with variations in the time-average capacity and not at a time slot rate,</u>" as required by claim 1 as amended. Thus, amended claim 1 is patentable over Uzun. Base claims 4 and 12 have been amended in a similar manner and are patentable for at least the same reasons as amended claim 1. Claims 2-3, 10, and 13-15, which depend from base claims 1, 4 and 12, are also allowable for at least the same reasons as amended claim 1. Reconsideration of the rejection under 35 USC 102 is respectfully requested.

To expedite prosecution of the application, claims 31 and 34 have been canceled. Applicants do not acquiesce to the rejection of claims 31 and 34 and reserve the right to pursue the subject matter of the canceled claims in a continuing application.

Claims 17-18 were rejected under 35 USC 102(e) as being anticipated by Miles, et al. (US6665495) ("Miles"). To expedite prosecution of the application, claims 17-18 have been canceled. Applicants do not acquiesce to the rejection of claims 17-18 and reserve the right to pursue the subject matter of the canceled claims in a continuing application.

Claims 5-9 and 33 were rejected under 35 USC 103(a) as being unpatentable over Uzun as applied to claims 4 and 31 and further in view of Krishna, et al. (US2001/0050916)("Krishna"). The rejection is respectfully traversed.

Claims 5-9 depend from base claim 4 and are allowable for at least the same reasons given for base claim 4. Reconsideration of the rejection under 35 USC 103 is respectfully requested.

To expedite prosecution of the application, claim 33 has been canceled. Applicants do not acquiesce to the rejection of claim 33 and reserve the right to pursue the subject matter of the canceled claim in a continuing application.

Claim 11 was rejected under 35 USC 103(a) as being unpatentable over Uzun as applied to claim 4 and further in view of Reeve, et al. (US2002/0027902)("Reeve"). The rejection is respectfully traversed.

Claim 11 depends from base claim 4 and is allowable for at least the same reasons given for base claim 4. Reconsideration of the rejection under 35 USC 103 is respectfully requested.

Claim 16 was rejected under 35 USC 103(a) as being unpatentable over Uzun as applied to claim 12 and further in view of Reeve and Miles. The rejection is respectfully traversed.

Claim 16 depends from base claim 12 and is allowable for at least the same reasons given for base claim 12. Reconsideration of the rejection under 35 USC 103 is respectfully requested.

Claims 19-27 and 29 were rejected under 35 USC 103(a) as being unpatentable over Uzun in view of Miles. The rejection is respectfully traversed.

Claim 19 has been amended to recite in part "a line card interface device for carrying traffic defined by a matrix with elements that specify the time-average capacity required between each traffic source – traffic sink pair, the capacity having a statistical distribution which varies at a given variation rate, ..., and the control device having means for applying control signals to the switch control inputs at a rate in accordance with the statistical variations in the time-average capacity and not at a time slot rate." As noted above, Uzun is directed to uniform traffic, i.e., traffic that has no statistical variations, and it reconfigures its central switching fabric once per time slot. Miles does not supply what is lacking in Uzun. Therefore, the combination of Uzun with Miles does not teach or suggest the elements of amended claim 19. Claims 20-27 and 29, which depend from base claim 19, are also allowable for at least the same reasons as amended claim 19. Reconsideration of the rejection under 35 USC 103 is respectfully requested.

Claim 28 was rejected under 35 USC 103(a) as being unpatentable over Uzun and Miles as applied to claim 21 and further in view of Carvey, et al. (US6934471)("Carvey"). The rejection is respectfully traversed.

Claim 28 depends from base claim 19 and is allowable for at least the same reasons given for base claim 19. Reconsideration of the rejection under 35 USC 103 is respectfully requested.

Claim 30 was rejected under 35 USC 103(a) as being unpatentable over Uzun in view of Horlin, et al. (US6993018). Claim 32 was rejected under 35 USC 103(a) as being unpatentable over Uzun as applied to claim 31 and further in view of Miles. To expedite prosecution of the

application, claims 30 and 32 have been canceled. Applicants do not acquiesce to the rejections of claims 30 and 32 and reserve the right to pursue the subject matter of the canceled claims in a continuing application.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By ______________________
Timothy J. Meagher
Registration No. 39,302
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Date: 8/4/9